

Limits of Breach-Resistant and Snapshot-Oblivious RAMs

Still $\Omega(b/w \cdot \log(nb/c))$


Giuseppe Persiano

Università di Salerno and Google LLC

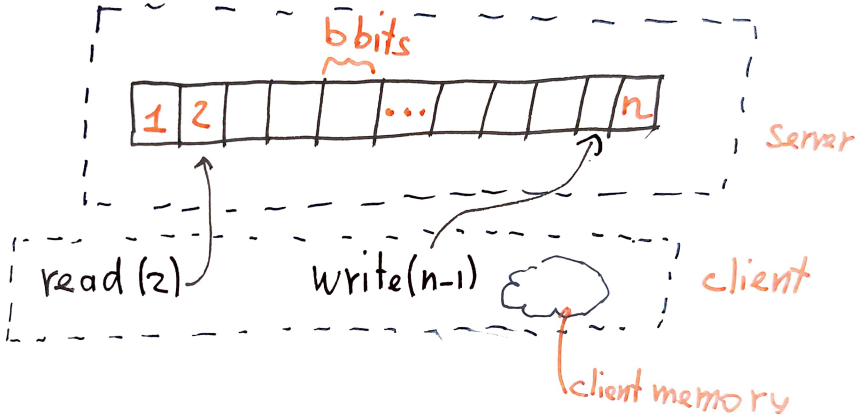
ESSA 3 in Bertinoro

Joint work with Kevin Yeo (Google LLC and Columbia U.)
Accepted to CRYPTO 23

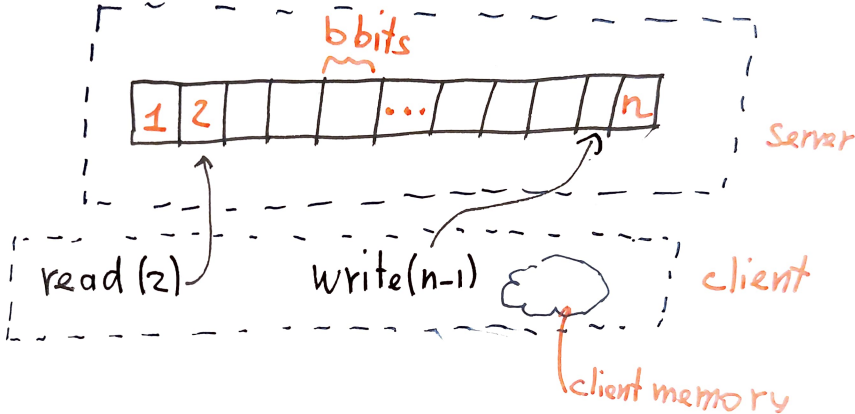
RAMS



RAMS



RAMS



Server memory words of length $w < b$

Oblivious RAMS

Hiding the access pattern to the RAM from the server

Upper bounds

- Goldreich Ostrovsky – Late 80's early 90'
 - ▶ Slowdown $O(\log^3 n)$
-
- Patel Persiano Raykova Yeo – 2018
 - ▶ Slowdown $O(\log n \log \log n)$
- Asharov Komogordsky Lin Peserico Shi – 2018
 - ▶ Slowdown $O(\log n)$

Oblivious RAMS

Hiding the access pattern to the RAM from the server

Upper bounds

- Goldreich Ostrovsky – Late 80's early 90'
 - ▶ Slowdown $O(\log^3 n)$
-
- Patel Persiano Raykova Yeo – 2018
 - ▶ Slowdown $O(\log n \log \log n)$
- Asharov Komogordsky Lin Peserico Shi – 2018
 - ▶ Slowdown $O(\log n)$

Lower bounds

- Larsen Nielsen – 2018
 - ▶ Slowdown $\Omega(\log n)$
- It holds also for Differential Privacy, some leakage

The snapshot adversary

the **Server** is the adversary

Snapshot Adversary

Du, Genkin, Grubbs, 2022

- The adversary gets control of the **Server** for L **consecutive** operations
 - ▶ Slowdown $O(\log L)$

The snapshot adversary

the **Server** is the adversary

Snapshot Adversary

Du, Genkin, Grubbs, 2022

- The adversary gets control of the **Server** for L **consecutive** operations
 - ▶ Slowdown $O(\log L)$

What if the adversary is active for more than one *window*?

Snapshot Resistant Stacks

- **read – write**

Snapshot Resistant Stacks

- **push – pop**

Snapshot Resistant Stacks

- **push – pop**
- want to hide sequence of operations
 - ▶ hide if **push** or **pop**
 - ▶ hide input to **push**

Snapshot Resistant Stacks

- **push – pop**
- want to hide sequence of operations
 - ▶ hide if **push** or **pop**
 - ▶ hide input to **push**
- **hiding from whom**
 - ▶ Adversary that can see S memory snapshots
 - ▶ Adversary that can see T operations transcripts

$(\infty, 0)$ -snapshot secure stacks


Adversary gets snapshots of memory after **all operations** and **no transcript**

$$S = \infty$$

$$T = 0$$



push(12)
push(24)
push(17)



$(\infty, 0)$ -snapshot secure stacks

Adversary gets snapshots of memory after **all operations** and **no transcript**

$$S = \infty$$

$$T = 0$$



push(17)
push(24)
push(17)

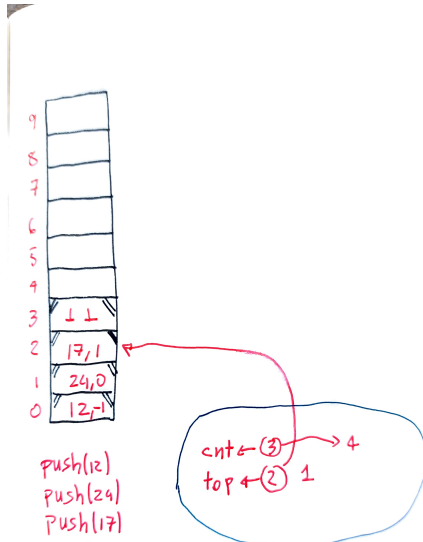
cnt ← 3
top ← 2

$(\infty, 0)$ -snapshot secure stacks

Adversary gets snapshots of memory after **all operations** and **no transcript**

$$S = \infty$$

$$T = 0$$



Snapshot Secure Stacks

- **Init()**

- ▶ randomly choose encryption key K
- ▶ set $\text{cnt} = 0$ and $\text{top} = -1$.

- **Push(v)**

- ▶ upload $\text{Enc}(K, (v, \text{top}))$ to location cnt
- ▶ set $\text{top} \leftarrow \text{cnt}$
- ▶ set $\text{cnt} \leftarrow \text{cnt} + 1$

- **Pop()**

- ▶ download pair (v, t) from location top
- ▶ upload a dummy encryption to location cnt
- ▶ set $\text{top} \leftarrow t$
- ▶ set $\text{cnt} \leftarrow \text{cnt} + 1$
- ▶ return v

Snapshot Security

Snapshots:

Only difference between operation i and operation $i + 1$ in location i

Snapshot Security

Snapshots:

Only difference between operation i and operation $i + 1$ in location i independent from history of ops

Snapshot Security

Snapshots:

Only difference between operation i and operation $i + 1$ in location i independent from history of ops

Transcripts:

Client reaches for the current `top`

Snapshot Security

Snapshots:

Only difference between operation i and operation $i + 1$ in location i independent from history of ops

Transcripts:

Client reaches for the current `top`

Gives information about number of `push` ops vs number of `pop` ops

Snapshot Security

Snapshots:

Only difference between operation i and operation $i + 1$ in location i independent from history of ops

Transcripts:

Client reaches for the current `top`

Gives information about number of `push` ops vs number of `pop` ops

randomly select a `PRP F` and write new pair at `F(cnt)`

$(\infty, 1)$ -snapshot secure stacks

\mathcal{A} gets snapshots of memory after every operations and transcript for one.

$(\infty, 1)$ -snapshot secure stacks

\mathcal{A} gets snapshots of memory after every operations and transcript for one.

Snapshot Secure Stacks

• Init()

- ▶ randomly choose seed S
- ▶ randomly choose encryption key K
- ▶ set $\text{cnt} = 0$ and $\text{top} = -1$.

• Push(v)

- ▶ download from location $F(S, \text{top})$ and discard
- ▶ upload $\text{Enc}(K, (v, \text{top}))$ to location $F(S, \text{cnt})$
- ▶ $\text{top} \leftarrow \text{cnt}$
- ▶ $\text{cnt} \leftarrow \text{cnt} + 1$

• Pop()

- ▶ download pair (v, t) from location $F(S, \text{top})$
- ▶ upload dummy encryption at location $F(S, \text{cnt})$
- ▶ set $\text{top} \leftarrow t$
- ▶ set $\text{cnt} \leftarrow \text{cnt} + 1$
- ▶ return v

This looks very promising

- Constant slowdown against *snapshot* adversary
 - ▶ *for the same price I can throw in one transcript of your choice*
- For *persistent* adversaries, stack is as hard as ORAM
 - ▶ Oblivious stack requires

$$\Omega(b/w)\log(nb/c)$$

Riko Jacob, Kasper Green Larsen, Jesper Buus Nielsen. SODA 2019.

This looks very promising

- Constant slowdown against *snapshot* adversary
 - ▶ *for the same price I can throw in one transcript of your choice*
- For *persistent* adversaries, stack is as hard as ORAM
 - ▶ Oblivious stack requires

$$\Omega(b/w)\log(nb/c)$$

Riko Jacob, Kasper Green Larsen, Jesper Buus Nielsen. SODA 2019.

- Actually, not! for ORAM is still

$$\Omega(b/w)\log(nb/c)$$

This looks very promising

- Constant slowdown against *snapshot* adversary
 - ▶ for the same price I can throw in one transcript of your choice
- For *persistent* adversaries, stack is as hard as ORAM
 - ▶ Oblivious stack requires

$$\Omega(b/w)\log(nb/c)$$

Riko Jacob, Kasper Green Larsen, Jesper Buus Nielsen. SODA 2019.

- Actually, not! for ORAM is still

$$\Omega(b/w)\log(nb/c)$$

Sorry

The snapshot adversary

Snapshot window (t, ℓ)

- A *snapshot window* of length ℓ starting at time t .
- The adversary receives
 - ▶ *snapshot* of server *memory content* before operation t has been executed
 - ▶ *transcript* of *server's operations* for the following ℓ operations that take place at times $t, t+1, \dots, t+\ell-1$.
- For $\ell = 0$, only memory content before operation t .

A (S, L) -snapshot adversary

Specifies a sequence of *snapshot windows* $\mathcal{S} = ((t_1, \ell_1), \dots, (t_s, \ell_s))$ such that

- $s \leq S$, at most S windows,
 - ▶ at most S snapshots
- $\sum \ell_i \leq L$, for a total duration of at most L operations
 - ▶ at most L transcripts

The Lower Bound

Theorem

For any $0 \leq \epsilon \leq 1/16$, let DS be a $(3, 1, \epsilon)$ -snapshot private RAM data structure for n entries each of b bits implemented over $w = \Omega(\log n)$ bits using client storage of c bits in the cell probe model. If DS has amortized write time t_w and expected amortized read time t_r with failure probability at most $1/3$, then

$$t_r + t_w = \Omega(b/w \cdot \log(nb/c)).$$

The Lower Bound

Theorem

For any $0 \leq \epsilon \leq 1/16$, let DS be a $(3, 1, \epsilon)$ -snapshot private RAM data structure for n entries each of b bits implemented over $w = \Omega(\log n)$ bits using client storage of c bits in the cell probe model. If DS has amortized write time t_w and expected amortized read time t_r with failure probability at most $1/3$, then

$$t_r + t_w = \Omega(b/w \cdot \log(nb/c)).$$

n *logical* blocks of b bits

The Lower Bound

Theorem

For any $0 \leq \epsilon \leq 1/16$, let DS be a $(3, 1, \epsilon)$ -snapshot private RAM data structure for n entries each of b bits implemented over $w = \Omega(\log n)$ bits using client storage of c bits in the cell probe model. If DS has amortized write time t_w and expected amortized read time t_r with failure probability at most $1/3$, then

$$t_r + t_w = \Omega(b/w \cdot \log(nb/c)).$$

$w < b$ is size *physical* words

The Lower Bound

Theorem

For any $0 \leq \epsilon \leq 1/16$, let DS be a $(3, 1, \epsilon)$ -snapshot private RAM data structure for n entries each of b bits implemented over $w = \Omega(\log n)$ bits using client storage of c bits in the cell probe model. If DS has amortized write time t_w and expected amortized read time t_r with failure probability at most $1/3$, then

$$t_r + t_w = \Omega(b/w \cdot \log(nb/c)).$$

Client has c bits of *local memory*

The Lower Bound

Theorem

For any $0 \leq \epsilon \leq 1/16$, let DS be a $(3, 1, \epsilon)$ -snapshot private RAM data structure for n entries each of b bits implemented over $w = \Omega(\log n)$ bits using client storage of c bits in the cell probe model. If DS has amortized write time t_w and expected amortized read time t_r with failure probability at most $1/3$, then

$$t_r + t_w = \Omega(b/w \cdot \log(nb/c)).$$

Adversary receives at most **3** memory *snapshots* and **1** operation *transcript*

The Lower Bound

Theorem

For any $0 \leq \epsilon \leq 1/16$, let DS be a $(3, 1, \epsilon)$ -snapshot private RAM data structure for n entries each of b bits implemented over $w = \Omega(\log n)$ bits using client storage of c bits in the cell probe model. If DS has amortized write time t_w and expected amortized read time t_r with failure probability at most $1/3$, then

$$t_r + t_w = \Omega(b/w \cdot \log(nb/c)).$$

ϵ is the adversary's advantage in the security game

The security game

$\text{Expt}_{\text{DS}, \mathcal{A}}^{n, \beta}$

- Receive $(O_0, O_1, ((t_1, \ell_1), \dots, (t_s, \ell_s)))$ from $\mathcal{A}_0(1^n)$.
- Set $\mathcal{L} \leftarrow \emptyset$, $\text{DS} \leftarrow (R_1, \dots, R_n)$, $i \leftarrow 1$.
- While $i \leq |O_\beta|$:
 - ▶ If $i = t_j$ for some $1 \leq j \leq s$:
 - ★ Set $\mathcal{L} \leftarrow \mathcal{L} \parallel (\text{memory}, M)$.
 - ★ For $k = 1, \dots, \ell_j$:
Execute $\text{DS}^{\text{LRead}, \text{LWrite}}(O_b[i])$ and set $i \leftarrow i + 1$.
 - ▶ Else:
Execute $\text{DS}^{\text{Read}, \text{Write}}(O_b[i])$ and set $i \leftarrow i + 1$.
- Return $\mathcal{A}_1(\mathcal{L})$.


$$\left| \Pr[\text{Expt}_{\text{DS}, \mathcal{A}}^{n, 0} = 1] - \Pr[\text{Expt}_{\text{DS}, \mathcal{A}}^{n, 1} = 1] \right| \leq \epsilon,$$

for all PPT \mathcal{A} that are (S, L) -snapshot adversaries.

The Epoch structure

The sequence and the epochs

- n logical indices
- $m \leftarrow \{n/2 + 1, \dots, n\}$
- m writes of **random** b -bit blocks at indices $1, 2, \dots, m$
- followed by **one** read.



A $(3, 1)$ -snapshot adversary – Intuition

- Two sequences of operations O_0, O_1

A $(3, 1)$ -snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices

A $(3, 1)$ -snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1

A $(3, 1)$ -snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch

A $(3, 1)$ -snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i

A $(3, 1)$ -snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**

A (3, 1)-snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**
 - ▶ epochs following epoch i are **not large enough**

A (3, 1)-snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**
 - ▶ epochs following epoch i are **not large enough**
 - ▶ pick i so that client memory is **too small**

A (3, 1)-snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**
 - ▶ epochs following epoch i are **not large enough**
 - ▶ pick i so that client memory is **too small**
- **correctness of O_0**
read of O_0 does not depend on epoch i

A (3, 1)-snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**
 - ▶ epochs following epoch i are **not large enough**
 - ▶ pick i so that client memory is **too small**
- **correctness of O_0**
read of O_0 does not depend on epoch i
- **security**
but if it does not, then security fails

A (3, 1)-snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**
 - ▶ epochs following epoch i are **not large enough**
 - ▶ pick i so that client memory is **too small**
- **correctness of O_0**
read of O_0 does not depend on epoch i
- **security**
but if it does not, then security fails
- **final step**
this holds for *all* epochs except for those that have fewer than c/b writes.

A (3, 1)-snapshot adversary – Intuition

- Two sequences of operations O_0, O_1
 - ▶ Both write **random** blocks to the first m indices
 - ▶ O_0 reads index 1
 - ▶ O_1 reads a randomly selected index j written in the i -th epoch
- **correctness of O_1**
touch about b/w cells updated in epoch i
 - ▶ epochs preceding epoch i are **independent**
 - ▶ epochs following epoch i are **not large enough**
 - ▶ pick i so that client memory is **too small**
- **correctness of O_0**
read of O_0 does not depend on epoch i
- **security**
but if it does not, then security fails
- **final step**
this holds for *all* epochs except for those that have fewer than c/b writes.
- **we have a lower bound $\Omega(b/w \cdot \log(nb/c))$**

A (3, 1)-snapshot adversary – Part 0

$\mathcal{A}_0^i(1^n)$

- Randomly select integer m from $[n/2, n]$.
- Randomly and ind. select $B_1, \dots, B_m \leftarrow \{0, 1\}^b$.
- Set $O_0 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(m))$.
- Randomly select $j \in [p_i, p_i + r^i - 1]$,
- Set $O_1 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(j))$.
- Set $\mathcal{S} = ((p_i, 0), (p_i + r^i, 0), (m + 1, 1))$.
- Return (O_0, O_1, \mathcal{S}) .

A (3, 1)-snapshot adversary – Part 0

$\mathcal{A}_0^i(1^n)$

- Randomly select integer m from $[n/2, n]$.
 - Randomly and ind. select $B_1, \dots, B_m \leftarrow \{0, 1\}^b$.
 - Set $O_0 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(m))$.
 - Randomly select $j \in [p_i, p_i + r^i - 1]$,
 - Set $O_1 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(j))$.
 - Set $S = ((p_i, 0), (p_i + r^i, 0), (m + 1, 1))$.
 - Return (O_0, O_1, S) .
-
- $(p_i, 0)$: **snapshot** of server memory before epoch i

A (3, 1)-snapshot adversary – Part 0

$\mathcal{A}_0^i(1^n)$

- Randomly select integer m from $[n/2, n]$.
- Randomly and ind. select $B_1, \dots, B_m \leftarrow \{0, 1\}^b$.
- Set $O_0 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(m))$.
- Randomly select $j \in [p_i, p_i + r^i - 1]$,
- Set $O_1 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(j))$.
- Set $S = ((p_i, 0), (p_i + r^i, 0), (m + 1, 1))$.
- Return (O_0, O_1, S) .

- $(p_i, 0)$: **snapshot** of server memory before epoch i
- $(p_i + r^i, 0)$: **snapshot** of server memory after epoch i

A (3, 1)-snapshot adversary – Part 0

$\mathcal{A}_0^i(1^n)$

- Randomly select integer m from $[n/2, n]$.
 - Randomly and ind. select $B_1, \dots, B_m \leftarrow \{0, 1\}^b$.
 - Set $O_0 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(m))$.
 - Randomly select $j \in [p_i, p_i + r^i - 1]$,
 - Set $O_1 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(j))$.
 - Set $S = ((p_i, 0), (p_i + r^i, 0), (m + 1, 1))$.
 - Return (O_0, O_1, S) .
-
- $(p_i, 0)$: **snapshot** of server memory before epoch i
 - $(p_i + r^i, 0)$: **snapshot** of server memory after epoch i
 - $(m + 1, 1)$: **snapshot** before read and **transcript** of read operation

A (3, 1)-snapshot adversary – Part 0

$\mathcal{A}_0^i(1^n)$

- Randomly select integer m from $[n/2, n]$.
- Randomly and ind. select $B_1, \dots, B_m \leftarrow \{0, 1\}^b$.
- Set $O_0 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(m))$.
- Randomly select $j \in [p_i, p_i + r^i - 1]$,
- Set $O_1 = (\text{write}(1, B_1), \dots, \text{write}(m, B_m), \text{read}(j))$.
- Set $S = ((p_i, 0), (p_i + r^i, 0), (m + 1, 1))$.
- Return (O_0, O_1, S) .

Important

- $(p_i, 0)$: **snapshot** of server memory before epoch i
- $(p_i + r^i, 0)$: **snapshot** of server memory after epoch i
- $(m + 1, 1)$: **snapshot** before read and **transcript** of read operation

A (3, 1)-snapshot adversary – Part 1

- U_i memory locations overwritten during epoch i
 - ▶ by comparing the **initial** and **final** snapshot of epoch i
- V_i memory locations overwritten since epoch i
 - ▶ by comparing the **final** snapshot of epoch i with snapshot **before the read**
- W_i memory location overwritten during epoch i that have not been modified when the read starts
 - ▶ $W_i = U_i \setminus V_i$
- Q_j cells from W_i read during **read(j)**,
- $|Q_j| \approx b/w$
 - ▶ \mathcal{A}^1 returns 0 iff $|Q_j| \leq \rho \cdot b/w$

The coding argument

Suppose

$$t_w = o(b/w \log(nb/c))$$

then there exists $\rho > 0$ such that, for most epochs i ,

$$|Q_j| \geq \rho \cdot b/w$$

with probability $\geq 1/8$ for j in epoch i .

The coding argument

Suppose

$$t_w = o(b/w \log(nb/c))$$

then there exists $\rho > 0$ such that, for most epochs i ,

$$|Q_j| \geq \rho \cdot b/w$$

with probability $\geq 1/8$ for j in epoch i .

Suppose not.

The coding argument

Suppose

$$t_w = o(b/w \log(nb/c))$$

then there exists $\rho > 0$ such that, for most epochs i ,

$$|Q_j| \geq \rho \cdot b/w$$

with probability $\geq 1/8$ for j in epoch i .

Suppose not.

Then we can encode the $r^i \cdot b$ bits of epoch i using fewer bits.

The coding argument - I

A coding game

- S wants to send B^i to R
 - ▶ the r^i blocks from epoch i
- S and R share
 - ▶ B^{-i} (all except epoch i)
 - ▶ randomness \mathcal{R} to execute DS.

$$\mathcal{H}(B^i | \mathcal{R}, B^{-i}) = r^i \cdot b.$$

The coding argument - II

- S and R execute all epochs $> i$

$\text{write}(1, B_1), \dots, \text{write}(p_i - 1, B_{p_i-1})$

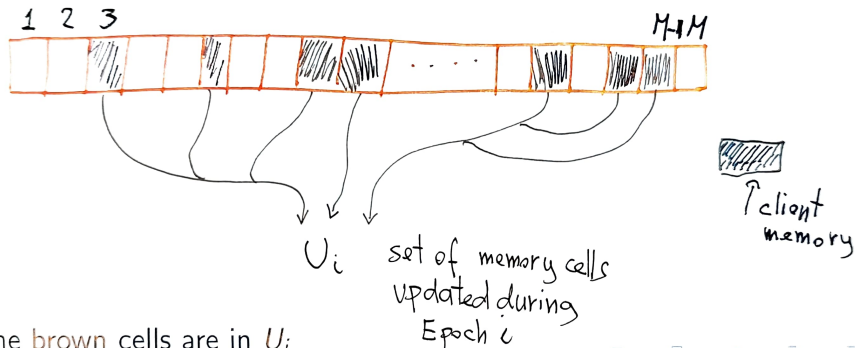


The coding argument

- S executes epoch i

$$\text{write}(p_i, B_{p_i}), \dots, \text{write}(p_i + r^i - 1, B_{p_i+r^i-1})$$


- Note: R cannot execute epoch i



The brown cells are in U_i


The coding argument

- S and R execute epochs $< i$
 - ▶ R needs some help
 - ★ client memory: c bits.
- For $j = p_{i-1}, \dots, m$
 - ▶ execute $\text{write}(j, B_j)$ touching T_j
 - ▶ R needs $U_i \cap T_j$ (cell location and content)



The coding argument

- c bits + set $Y_i := U_i \cap (T_{p_i+r_i} \cup \dots \cup T_m)$



The coding argument

\mathbb{S} memory state after $\text{write}(m, B_m)$

- For $j = p_i, \dots, p_i + r^i - 1$
 - ▶ S and R execute $\text{read}(j)$ starting from \mathbb{S}
 - ▶ R needs $Q_j := W_i \cap T_j^m$
 - ▶ if read errs or $Q_j > \rho b/w$
 - ★ B_j is added to encoding
 - ▶ else
 - ★ Q_j is added to encoding

Length of encoding

Length depends on

- Set Y_i
 - ▶ for most epochs i , $\mathbb{E}[|Y_i|] \leq r^{i-1}b/w$
- Set Q_j
 - ▶ By assumption $|Q_j| < \rho \cdot b/w$ with prob $\geq 7/8$

Length of encoding

Length depends on

- Set Y_i
 - ▶ for most epochs i , $\mathbb{E}[|Y_i|] \leq r^{i-1}b/w$
- Set Q_j
 - ▶ By assumption $|Q_j| < \rho \cdot b/w$ with prob $\geq 7/8$

Encoding is too small

Getting there...

$$t_w = o(b/w \log(nb/c))$$

implies that, for most epochs i ,

$$|Q_j| \geq \rho \cdot b/w$$

with probability $\geq 1/8$ for j in epoch i . from epoch i .

Getting there...

$$t_w = o(b/w \log(nb/c))$$

implies that, for most epochs i ,

\mathcal{A} outputs 1 with probability $\geq 1/8$ when reading j from epoch i .

Getting there...

$$t_w = o(b/w \log(nb/c))$$

implies that, for most epochs i ,

\mathcal{A} outputs 1 with probability $\geq 1/8$ when reading j from epoch i .

If $\epsilon = 1/16$ then \mathcal{A} outputs 1 with probability $\geq 1/16$ when reading m

Getting there...

$$t_w = o(b/w \log(nb/c))$$

implies that, for most epochs i ,

\mathcal{A} outputs 1 with probability $\geq 1/8$ when reading j from epoch i .

If $\epsilon = 1/16$ then \mathcal{A} outputs 1 with probability $\geq 1/16$ when reading m

$\text{read}(m)$ must touch $\geq \rho \cdot b/w$ cells from epoch i

Getting there...

$$t_w = o(b/w \log(nb/c))$$

implies that, for most epochs i ,

\mathcal{A} outputs 1 with probability $\geq 1/8$ when reading j from epoch i .

If $\epsilon = 1/16$ then \mathcal{A} outputs 1 with probability $\geq 1/16$ when reading m

$\text{read}(m)$ must touch $\geq \rho \cdot b/w$ cells from epoch i

$$\Omega(b/w \cdot \log nb/c)$$

Wrapping up

Now...

If writes are fast

$$t_w = o(b/w \log(nb/c))$$

then $\text{read}(j)$ in epoch i has $Q_j = \Omega(b/w)$ with prob at least $1/8$.

Wrapping up

Now...

If writes are fast

$$t_w = o(b/w \log(nb/c))$$

then $\text{read}(j)$ in epoch i has $Q_j = \Omega(b/w)$ with prob at least $1/8$.

Reading 1

Must touch from each large epoch $O(b/w)$ cells otherwise we lose security.

$$\Omega(b/w \cdot \log(nb/c))$$